

(NASA-CR-194876) SUPERCOMPUTER  
NETWORKING FOR SPACE SCIENCE  
APPLICATIONS Final Report, 29 Jan.  
- 31 Dec. 1991 (George Washington  
Univ.) 42 p

N94-23085

Unclass

G3/62 0203553



## Report Documentation Page

11-62-12

1. Report No.  Final		2. Government Accession No.		3. Recipient's Catalog No.  203453 42P	
4. Title and Subtitle  Supercomputer Networking for Space Science Applications				5. Report Date  May 15, 1992	
				6. Performing Organization Code  IASR	
7. Author(s)  Professor B. I. Edelson				8. Performing Organization Report No.	
				10. Work Unit No.  506-59-11	
9. Performing Organization Name and Address  Institute for Applied Space Research George Washington University Washington, DC 20052				11. Contract or Grant No.  NAS5-30428	
				13. Type of Report and Period Covered  Final 1/29-12/31/91	
12. Sponsoring Agency Name and Address  CESDIS Code 930.5 Goddard Space Flight Center Greenbelt, MD 20771				14. Sponsoring Agency Code	
15. Supplementary Notes  CESDIS, the Center of Excellence in Space Data and Information Sciences) is operated by Universities Space Research Association (USRA), American City Building, Suite 212, 10227 Wincopin Circle, Columbia, MD 21044.					
16. Abstract  This report covers the initial design of a supercomputer network topology including the design of the communications nodes along with the communications interface hardware and software. The report also characterizes several space science applications that are proposed experiments by GSFC and JPL for a supercomputer network using the NASA ACTS satellite.					
17. Key Words (Suggested by Author(s))  supercomputer ACTS				18. Distribution Statement  unclassified, unlimited	
19. Security Classif. (of this report)		20. Security Classif. (of this page)		21. No. of pages  37	
				22. Price	

## Table of Contents

1. Executive Summary
2. Introduction
3. Network Topology
  - 3.1 Wide Area Network Topology
  - 3.2 Local Area Network Topology
4. Design of the Point-to-Point Communications Link
  - 4.1 Host-to-Host and Host-to-FEP Communications
  - 4.2 Spacelink Communications
  - 4.3 ACTS Earth Terminal Requirements
  - 4.4 ACTS Earth Terminal Description
5. ACTS Experiment Applications
  - 5.1 Science Data Analysis Applications
    - 5.1.1 Hubble Spacecraft Data Analysis
    - 5.1.2 Computational Fluid Dynamics Simulation
    - 5.1.3 Atmospheric Temperature Calculations
    - 5.1.4 TOVS Pathfinder
    - 5.1.5 Coupled Atmosphere-Oceans Climate Model
  - 5.2 Animation and Visualization of Movie Frames
  - 5.3 Expansion of Potential Applications
6. Conclusions and Recommendations
7. List of Investigators and Contributors
8. Bibliography

## **1. Executive Summary**

NASA has a long tradition, beginning in the early 1960s with such projects as Echo, Relay, Syncom and the Applied Technology Satellite (ATS) program, of developing advanced technologies for satellite communications. The Advanced Communications Technology Satellite (ACTS) system, scheduled to be operational in early 1993, continues this tradition providing for the first time very high capacity transmission circuits, with on-board processing. These capabilities allow the ACTS system to play a major role in the establishment of high data rate networks that are suitable for the interconnection of supercomputers, data archives, workstations and visualization devices into a Distributed High Performance Computing (DHPC) environment.

The ACTS system should demonstrate convincingly that communications satellites, in conjunction with terrestrial optical fibers, must be integral parts of our nation's High Performance Computing and Communications (HPCC) initiative.

To demonstrate the data distribution capabilities of communications satellites in general and the ACTS system in particular, a team consisting of The George Washington University, Comsat Laboratories, Cray Research, Inc., the NASA Goddard Space Flight Center (GSFC) and the NASA Jet Propulsion Laboratories (JPL), has proposed an ACTS experiment entitled "Supercomputer Networking Applications." That proposal has received initial acceptance as an ACTS experiment.

The study described in this document provides the initial design of the topology of a supercomputer network intended to interconnect a variety of high speed data processing, storage and display devices into a distributed computational facility. It also focuses attention on several potential Earth and space science applications experiments to be carried out over the distributed computational facility. A unique aspect of this networking experiment involves the interconnection of the supercomputers of the participating NASA centers, via the ACTS system, and the ultimate extension of the network to supercomputer centers that are currently part of the National Research and Educational Network (NREN) and/or the current terrestrial gigabit testbeds. Thus, the overall network

is configured to incorporate interactive satellite and terrestrial high data rate computing and communications, working together to solve real time science applications.

In this report we present the final results of our study specific to the link connecting supercomputers at the NASA Goddard Space Flight Center and the Jet Propulsion Laboratory. The network design concentrates on the architectural components of the network and includes the specification of the point-to-point communications link provided by the ACTS satellite, together with the requirements of the communications interface hardware and software that allow the connection of the computing equipment to the satellite terminals. The design also includes the specification of the ground terminal equipment that transmits the data to the ACTS satellite.

In addition to the technical specifications, the report provides a description of several actual applications that are planned to be carried out over the distributed computing environment. Joint discussions by GSFC and JPL personnel have revealed a large number of significant applications, six of which are described in this study. These are primarily space science applications that reside in the archives of one of the NASA centers and normally involve scientific analysis by investigators at both center locations, or at remote entry terminals. The experiments will utilize Parallel Virtual Machine (PVM) software that allows the distribution of an application over geographically disbursed computing resources, and three-dimensional data visualization equipment for information display.

It is planned to begin with simple computational data analysis tasks, progressing to more sophisticated applications such as a "Coupled Atmosphere-Oceans Climate Model". In this dynamic modeling application, GSFC is responsible for providing the atmosphere data, with JPL offering the oceans data. Parallel processing and dual manipulation of the data are essential ingredients in the solution of the problem. This parallel assignment experiment is intended as a first step toward the development of the Distributed Active Archive Centers of the EOSDIS program and will allow the demonstration of the ability to interconnect and distribute concurrent computing resources to geographically remote locations via the ACTS system. In subsequent experimentation the climate model will be coupled with research efforts at other locations

to distribute a similar model across multiple supercomputers linked by one of the regional terrestrial gigabit testbeds.

The conclusions of our design study reflect the important dimensions of supercomputer networking for the ACTS program. The proposed applications have not been performed in previous programs, nor can they be performed in real time with current communications satellites. This importance is also being recognized by other federal research centers, universities, and industrial organizations that have supercomputers or high data rate requirements. Our team is being approached by a number of these entities with requests to join our experiment or to receive reports or documentation that result from our experiment.

This study confirms that supercomputer networks using the ACTS satellite are feasible at bit rates up to 800 Mbit/s. Networking can be accomplished with either FDMA, TDMA or CDMA techniques to create multiple communications channels over single transponders. An examination of data rates concludes that the OC-3 rate of 155 Mbit/s can be implemented, and be ready for tests and demonstrations shortly after the launch of the satellite. In addition, the OC-3 rate is compatible with terrestrial networks that are planned to be available for testing during the same period.

The necessary earth terminal and interface equipment, described in this report, are presently not available. A joint DARPA/NASA procurement (BAA92-12) has been issued to acquire the ground terminals, but it is not clear whether the terrestrial interface units (TIU), and the front-end processor (FEP) hardware and software are included in the procurement. Clearly this interface equipment is essential in order to connect a supercomputer to the ground terminal equipment.

The study notes and emphasizes the urgent need to procure the necessary terminal and corresponding interface equipment for the experiment. If the interface equipment is not covered in BAA92-12, a special effort is required to complete the specifications and procure the TIU and FEP hardware and software.

The study recommends that the importance of supercomputer networking be recognized with terminals and experiment opportunities provided not only to our GSFC and JPL effort, but to other NASA centers such as LeRC, ARC, and LaRC, other government

research centers such as the Naval Research Laboratory, the DOE National Laboratories, and the National Center for Atmospheric Research. In addition, the University of Hawaii has established a budget for the procurement of a high data rate terminal so it can be a part of the ACTS experimental program.

Finally, we point out the importance of demonstrating the capability of ACTS to become an integral part of the nation's Grand Challenges: High Performance Computing and Communications (HPCC) network. While the HPCC backbone system will be based on terrestrial fiber optic communications facilities, there will always be special situations, such as remote or temporary locations, redundancy requirements and other valid reasons for communications satellites to be active participants in the network. Therefore, the national communications carriers, which will provide facilities for the HPCC network, should be actively encouraged to participate in the ACTS DHPC experiments.

## **2. Introduction**

In the past 15 years Distributed High Performance Computing (DHPC), as a means to solve large scale scientific and engineering problems, has become an essential tool in many areas of NASA's research and development program. In its usual form, DHPC involves the decomposition of a problem solution into several subtasks, their assignment to individual computing resources and the interaction of the geographically distributed computing resources, data archives and visualization devices in order to achieve a speedup of the time required to complete the solution of the problem. Among its many advantages the following are of particular importance:

- DHPC brings to bear the computational power of more than one machine and uses the most suitable machine for a specific subtask
- DHPC can access centrally located databases that may be too costly or too impractical to replicate
- DHPC can provide a local visualization of results produced by remote applications programs.

Within NASA, DHPC finds its most important applications in aeronautics, Earth and space sciences and space exploration, where sophisticated computer models and simulations of devices and processes and the visualization of large volumes of scientific data have become the primary means of analysis, synthesis and design. Many scientific and engineering problems lend themselves to analysis in a DHPC environment. A representative but by no means exhaustive list includes the following:

- Modeling and simulation of aerodynamic processes
- Modeling of coupled atmosphere-ocean processes for global climate studies
- Modeling of quantum chemical reaction dynamics in combustion processes, atmospheric chemistry and plasmas
- Three-dimensional modeling of the earth's crust for seismic profiling



- Interactive three-dimensional rendering of multiple earth science data sets and data visualization at remote workstations
- Supercomputer studies of global change among remote research sites
- Solution of real time very large data set problems partitioned over several supercomputers
- Sharing of the computational workload among several supercomputers by dividing the work into fine grain and coarse grain aspects and performing processing on the most appropriate machine

An important part of the emerging DHPC environment is the development of high data rate communications networks that can be used to interconnect the geographically distributed databases, computational resources and visualization devices, and to provide remote access to these resources. Several research projects, supported by both the United States Government and private enterprise, have recently been initiated to specify the requirements, develop the necessary communications technology and eventually implement networks that would offer data rates in the gigabit per second range on a national basis. Perhaps the most important of these is the gigabit testbed project being pursued by the Advanced Research Projects Agency (DARPA) of the Department of Defense.

While most of the emphasis in the development of high speed networking has been on the use of optical fibers as the communications medium, NASA has recently started to explore the possibility of employing geostationary communications satellites to provide the high data rate communications channels required for the interconnection of its DHPC facilities. Satellites enjoy certain advantages over terrestrial circuits and can play an important role in the total high speed data communications infrastructure. Three major areas of application taking particular advantage of the satellite's capabilities can be identified.

First, since the transmission capacity of a satellite system can be brought to bear anywhere within its coverage area at a cost that is largely independent of location, the system can economically offer data communications on a global basis, even to areas that

are traditionally underserved by terrestrial networks. Thus, remote areas of the world or areas in which terrestrial services cannot be justified economically due to limited demand may benefit from the availability of satellite coverage.

Second, the satellite system can serve as a backup for terrestrial data communications networks during periods of time when the latter experience service interruptions due to physical failures or other causes. This capability may result in major increases in the availability of telecommunications services offered by the terrestrial networks.

Third, the circuits of a satellite system can be used as standby channel capacity to manage traffic overloads that may on occasion arise over the terrestrial networks. The availability of such capacity allows terrestrial networks to be designed more closely in line with average rather than peak loads and may result in major cost savings.

Perhaps the most important advantage from the point of view of NASA's immediate needs is the near term availability of high data rate satellite communications channels over the Advanced Communications Technology Satellite (ACTS). Whereas the deployment of fiber may occur over a time frame of a decade or more, the ACTS system will offer high speed communications between NASA centers and other DHPC installations shortly after launch, which is now scheduled for early 1993. Through the technology of hopping spot beams and the on-board microwave switch matrix its channel capacity, which can be as high as several hundred Mbit/s, can be brought to bear anywhere within the Continental United States on demand. ACTS can therefore be expected to play an important part in the interconnection of NASA's supercomputers, data archives and visualization laboratories and thus fill a vital need in NASA's research and development programs.

NASA is currently in the process of specifying several high speed networking experiments utilizing the unique features of ACTS. We list several of the more important applications:

- Interconnection of the terrestrial gigabit network testbeds via high speed ACTS channels established through the microwave switch

- Downloading of data produced by applications programs running on supercomputers to remote workstations
- Extension of operating systems across ACTS, to allow loosely coupled applications programs running on Cray machines and utilizing distributed databases and subroutines in an open distributed environment
- Remote access by multiple users to supercomputer resources located at the Goddard Space Flight Center, the Jet Propulsion Laboratory and other NASA centers
- Tandem connections of supercomputer resources at GSFC, JPL and the San Diego Supercomputer Center via ACTS channels and terrestrial fiber circuits.

In support of NASA's efforts toward the implementation of these experiments The George Washington University, in close cooperation with Comsat Laboratories, Cray Research, the Goddard Space Flight Center and the Jet Propulsion Laboratory, has over the last 9 months performed a number of tasks under USRA Subcontract 550-69.

Generally, the objective of this work is to contribute to the development of the concepts, plans and techniques for the utilization of ACTS in NASA's high-speed distributed computing environment. Specifically, the goal is to design an experiment through which the capability of ACTS to provide high quality communications services in the high performance distributed computing environment can be demonstrated.

As a concrete project we focus on the communications services provided by ACTS in the implementation of real time distributed simulations and three dimensional visualization of processes such as shock waves, vortices, shear layers and wakes that represent atmospheric processes. To carry out the computational and visualization tasks requires the exchange of high volumes of data between a pair of supercomputers and several data archives over the ACTS communications channels. The experiment is designed to demonstrate the capabilities of the ACTS system in providing the necessary connectivity, high data throughput, low delay and fast channel access among the supercomputers, data archives and display devices.

This document is the final report under Subcontract 550-69. It contains the specifications for the design of the physical communications link, the terrestrial interface equipment and the hardware and software requirements for the applications experiment. Also included are descriptions of a number of candidate applications experiments for eventual implementation over the supercomputer network.

Section 2 discusses the topological aspects of a general wide area network interconnecting the geographically distributed resources of the DHPC environment. It also investigates the characteristics of the local area networks that provide the interconnection of the local resources of a supercomputer center, visualization laboratory or data archive.

Section 3 considers the architectural aspects of the physical and logical interactions between the various components of the point-to-point networking configuration.

Section 4 focuses on the design of the communications link. It explores the major issues of the terrestrial interface between ACTS and the local area networks and provides a consideration of the required modifications to the standard protocols to accommodate the large delay and high data rate of the ACTS. This section also provides a detailed specification of the protocol reference model that governs the logical and physical interactions between the various components of the end-to-end system.

Section 5 specifies the characteristics of the application to be implemented over the DHPC system and provides a listing of the planning and scheduling tasks to carry out the experiment.

Section 6 offers a summary of our conclusions and provides recommendations for specific implementation steps.

Section 7 is a list of the investigators and contributors to this study and to the planned ACTS Supercomputer Networking for Space Science Applications experiments.

Section 8 is a brief bibliography of the most recent technical documents on high data networking via the ACTS satellite, as well as documents relating to applications and experiments mentioned in this report.

### **3. Network Topology**

#### **3.1 Wide Area Network Topology - ACTS and the CASA Network**

The supercomputers to be interconnected as part of the ACTS networking experiment are currently located at the Jet Propulsion Laboratories (JPL), the California Institute of Technology (CIT), the San Diego Supercomputer Center (SDSC), the Los Alamos National Laboratory (LANL) and the Goddard Space Flight Center (GSFC). The network may also include the Cray computer at the NASA Ames Research Center. The specific supercomputer configurations include:

- Goddard Space Flight Center  
Cray Model Y-MP
- Jet Propulsion Laboratory  
Cray model X-MP/18 and its upgrade to a Cray Model Y-MP2E
- California Institute of Technology  
Hypercube Mark IIIfp
- San Diego Supercomputer Center  
Cray Model Y-MP/832
- Los Alamos National Laboratory  
CM-2 and Cray Model X-MP/416

A part of the network will be based on terrestrial common carrier provided gigabit fiber optic links. These will be installed from JPL to CIT, from CIT to SDSC, and from SDSC to LANL as part of the CASA gigabit testbed development. Initially only the interconnection of the JPL and GSFC machines will be based on ACTS communications channels.

Figure 1 shows the initial topological configuration of the terrestrial part of this network.

# CASA GIGABIT WAN

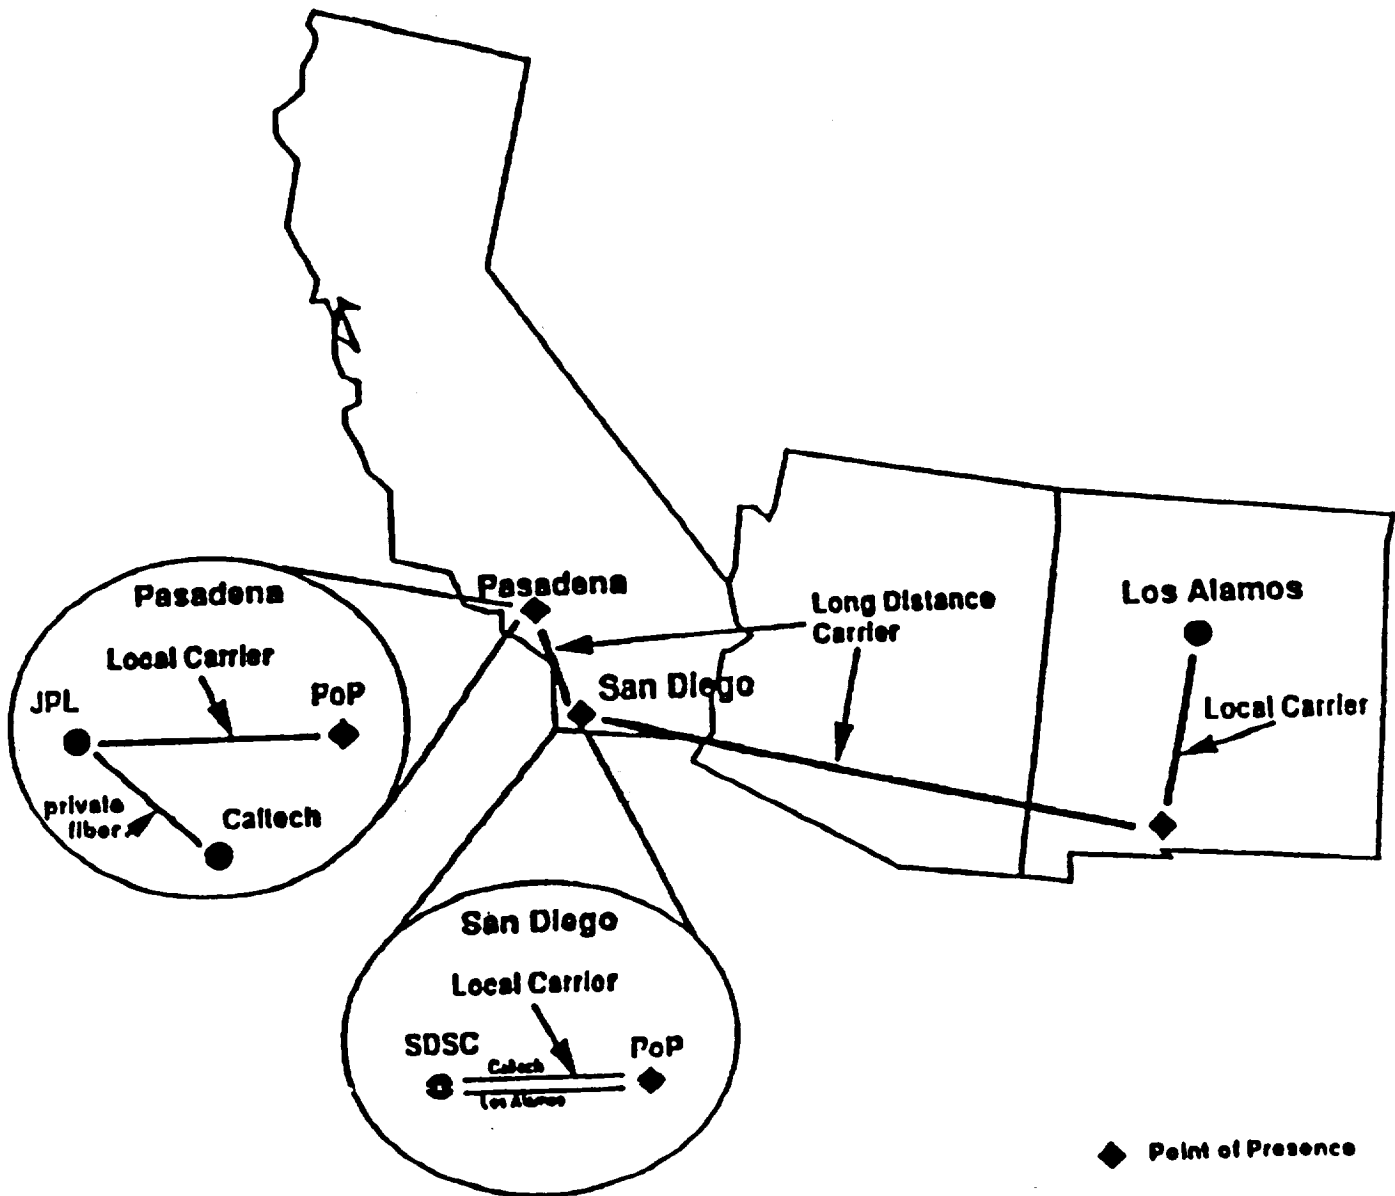


Figure 1

Topology of Distributed High Performance Computing Network

The configuration of the point-to-point ACTS link between the GSFC and JPL supercomputers is shown in Figure 2. The link consists of a pair of supercomputers that are interconnected via front-end-processors, terrestrial interfaces and the high data rate terminal. The latter includes the forward error correction codec, a modem, a high power amplifier and additional ground terminal transmitting and receiving equipment.

The details of the terrestrial part of the system are shown in Figure 3. Included here are the host computer, a communications front-end processor and the equipment necessary to convert the signals and protocols from terrestrial formats to those required on the space link.




Figure 2  
Point-to-Point Communications Link



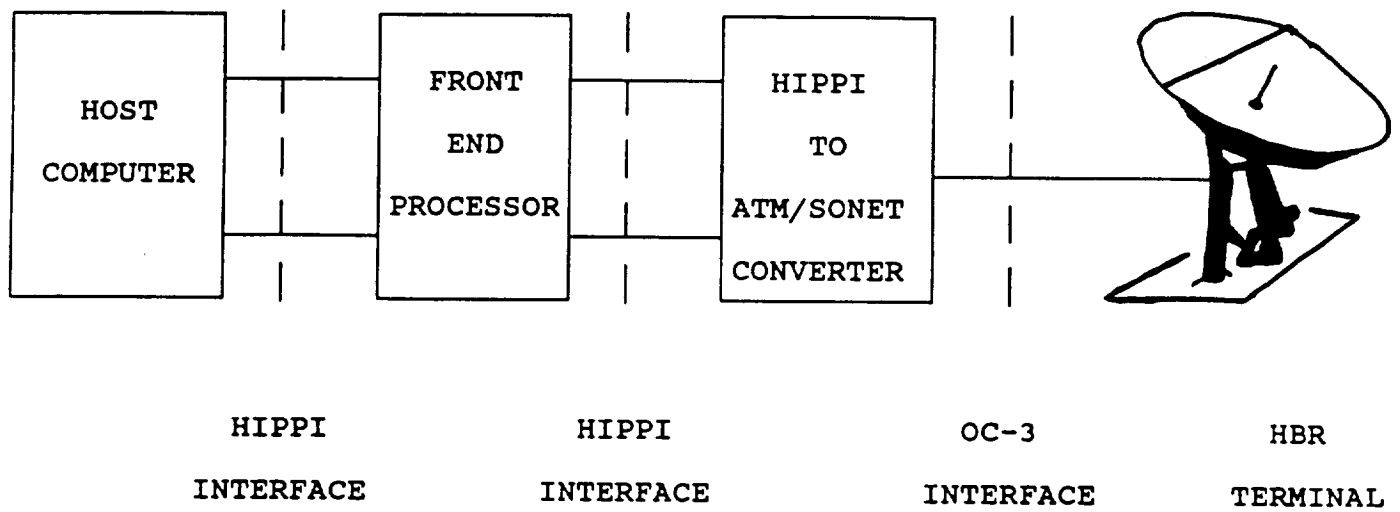


Figure 3  
Terrestrial Communications Interface

## **3.2 Local Area Network Topology**

### **1. NASA Goddard Space Flight Center**

Local communications between the various devices that make up the supercomputer installation at the GSFC are provided by an Ultranet Ring Network. This structure connects supercomputers, file servers and workstations distributed over several square miles of the GSFC campus. The Ultranet runs at gigabit per second rates and is the first such high speed ring in existence. Interfacing of the satellite link with the Ultranet occurs at the Ultranet 1000 hub. The NASA Goddard supercomputer and ultranet configuration are shown in Figure 4.

# NASA GODDARD SPACE FLIGHT CENTER

## Bldg. 28 - Supercomputer/Information Center




Figure 4  
GSFC Ultranet Configuration

## 2. NASA Jet Propulsion Laboratory

The Jet Propulsion Laboratory/California Institute of Technology computer network has an Ultranet HSC/HIPPI 800 Mbit/s capability that interconnects the supercomputers at the JPL/CIT campus over a dedicated fiber optic link. This JPL/CIT computer network will be connected to the CASA regional gigabit testbed facility. The network, with its FDDI link and local area Ethernet link, is shown in Figure 5.

**JPL**

# JPL/CIT COMPUTER NETWORK

CALTECH

JPL  
CASA-Net

UltraNet and/or HSC: 800Mb/s ('90-'91)

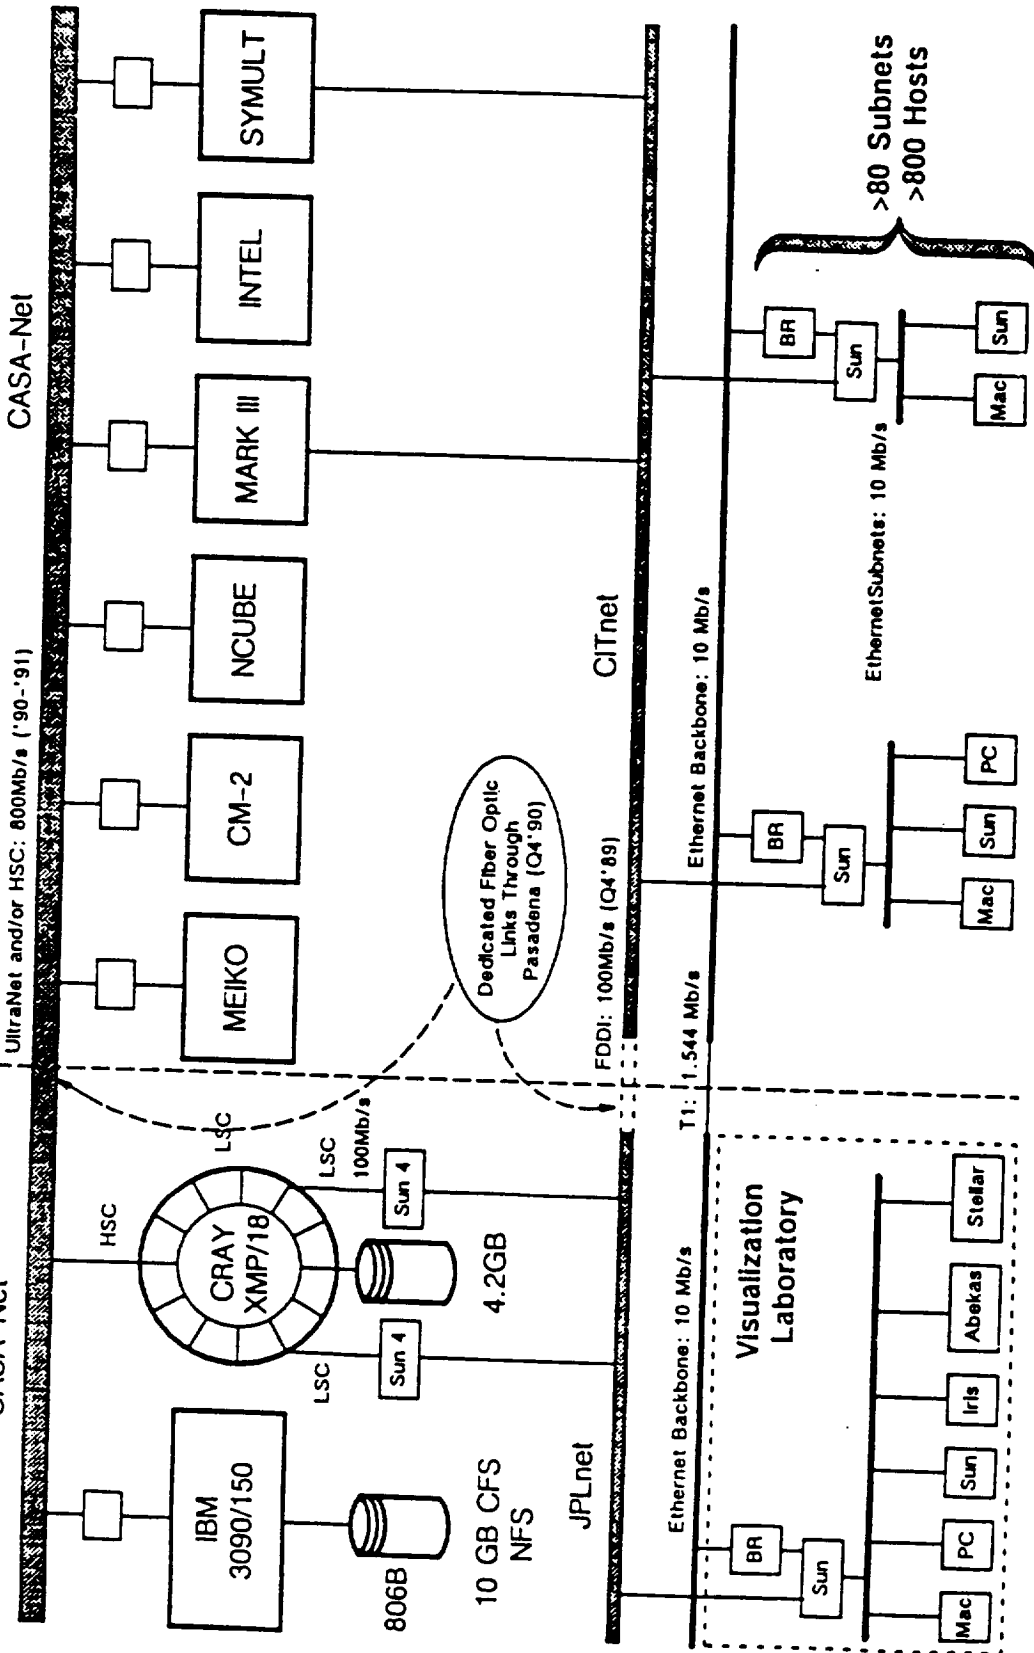


Figure 5

JPL Local Area Network Installation

#### **4. Design of the Point-to-Point Communications Link**

In this section we explore the issues involved in the provision of a point-to-point link via the ACTS system between a pair of supercomputers located at the Jet Propulsion Laboratory and the Goddard Space Flight Center.

##### **4.1 Host-to-Host and Host-to-FEP Communications**

The protocol reference model governing the end-to-end interactions between the host computers at JPL and GSFC, as well as the relationship between a host computer and its associated front-end processor is shown in Figure 6. The model is based on the well known Open Systems Interconnection Reference Model (OSIRM) of the International Organizations for Standardization (ISO).

As shown in the model, above layer 3 the communication between host computers is end-to-end and therefore does not involve the satellite system. It utilizes standard software implementations of Unix, Express and Strand. Layer 7 employs various applications packages typified by the Calcrust, Landsat, SAR and other rendering applications.

At layer 3 and below the host interacts only with its front-end processor. At layer 1 communication takes place over an 800 Mbit/s HIPPI connection. Layer 2 conforms to the standard layer 2 HIPPI protocol and performs frame synchronization, error control and flow control. For CRAY X-MP or Y-MP with Model C or D Input-Output System (IOS), the HIPPI socket is on XIOP, with one HIPPI channel per XIOP. For CRAY Y-MP2E, the HIPPI interface is via the HCA-3 and HCA-4 adapters. In the CRAY X-MP and Y-MP computers the HIPPI interface is supported by UNICOS 6.0 software.

All layer 3 functions between host and front-end processor are performed by a modified version of the IP protocol.

The purpose of the front-end processor is to provide decoupling of the host computer communications functions and its processing functions by allowing the host to initiate data transfer at any time, without regard to the availability of the communications




Figure 6

Host-to-Host and Host-to-FEP

Protocol Reference Model

link, the status of the error control procedure or flow control. The front-end processor also provides smoothing of the data flow over the communications link and allows for traffic scheduling and host interrupt.

Specific functions performed by the front-end processor include the following:

- Buffering of transmitted and received data to be processed by the host computer
- Layer 2 protocol processing, including framing, encapsulation, and address translation for routing

Several front-end processors are currently under development. Carnegie Mellon University, Intel, Bellcore and Network Systems Corporation are developing the Communications Accelerator Board (CAB) as part of the NECTAR gigabit testbed. The first version is currently in the design phase, with the prototype expected in the fall of 1992. The CAB supports HIPPI interfaces, a Unix socket and TCP/IP. Also in the design stage is the Crossbar Interface (CBI) being developed by Los Alamos National Laboratory as part of the CASA gigabit testbed. The CBI operates at multiples of 155 Mbit/s. Its prototype is expected to be available by the end of 1992.

## **4.2 Spacelink Communications**

Turning now to the spacelink, Figure 7 depicts the protocol reference model for the interconnection of two front-end processors via a signal converter, an ACTS High Data Rate (HDR) terminal and the ACTS-provided space communications channel.

The converter shown in Figure 7 provides protocol conversion between the HIPPI formats that govern the relationship between Host and FEP and the ATM/SONET transmission formats on the spacelink. A functional block diagram of the converter is shown in Figure 8 and the corresponding protocol reference model is shown in Figure 9.






Figure 7

FEP-to-FEP Protocol Reference Model

----- HIPPI Interface

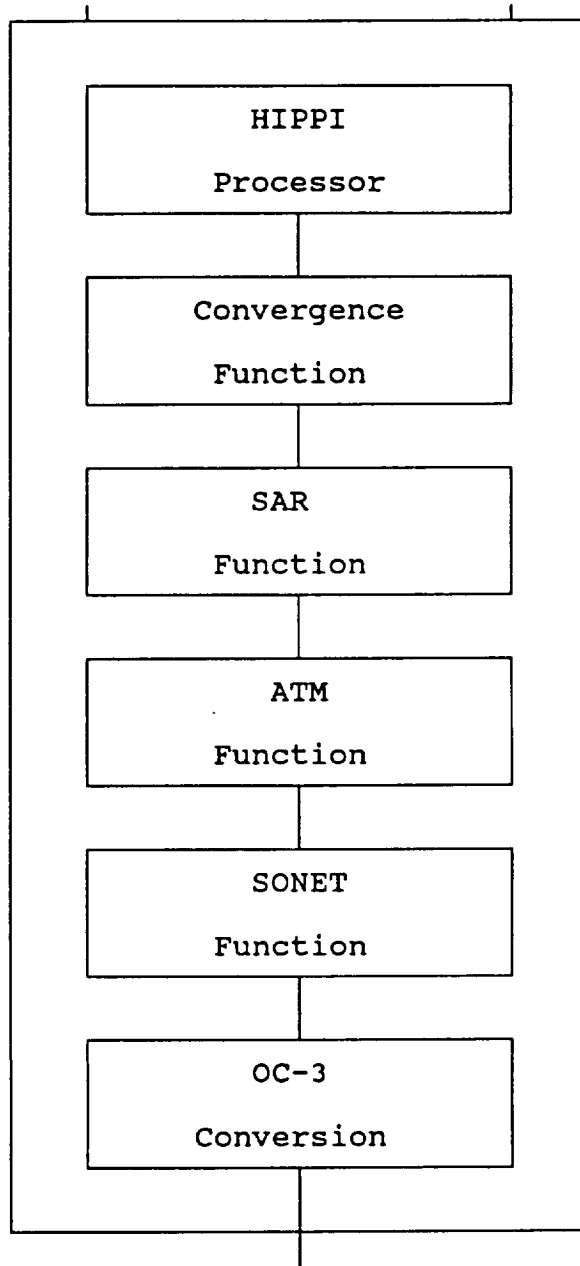
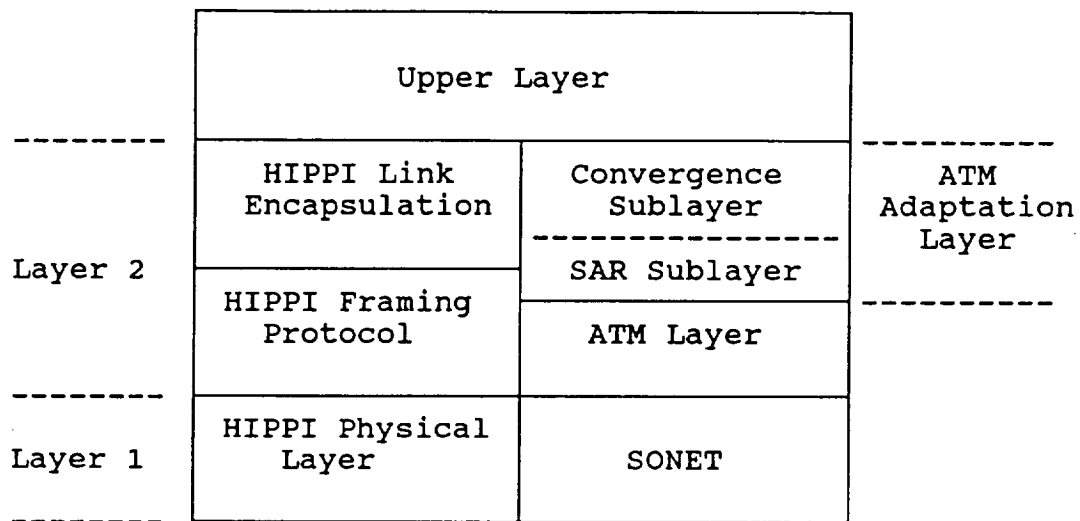


Figure 8

HIPPI-to-ATM/SONET Converter



HIPPI: High Performance Parallel Interface

SAR: Segmentation and Reassembly

SONET: Synchronous Optical Network

ATM: Asynchronous Transfer Mode

Figure 9

HIPPI-to-ATM/SONET Converter  
Protocol Reference Model

The converter provides the following functions:

- Full duplex transmission and reception of HIPPI PH-Bursts at 800 Mbit/s over separate HIPPI channels
- Conversion between HIPPI PH-Bursts, FP-Packets, HIPPI LE-Packets, upper layer data units and Convergence Packets, Segmentation and Reassembly Segments and ATM Cells
- Multiplexing of ATM Cells into SONET frames, SONET framing, transmission and reception at 155 Mbit/s (STS-3)
- Conversion to and from electrical to optical signals and formats (OC-3)

The converter performs only SONET processing, but no higher layer protocol processing. The specific layer functions are given in Table 1.

A prototype of the HIPPI/ATM/SONET converter is currently under development by Bellcore and Carnegie Mellon as part of the NECTAR gigabit testbed. Availability is projected to be the end of 1992.

Next we consider the major communications parameters of the space link.

- Data Rate

It is generally agreed that the experiments described in Section 5 below require the interconnection of supercomputers, data archives and visualization equipment at data rates in the several hundred megabits per second range. For purposes of compatibility with future terrestrial telecommunications capabilities the STS-3/OC-3 rate of 155.52 Mbit/s was selected.

- Bit Error Rate

For efficient utilization of the processing power of supercomputers, and to minimize the information transfer delay, bit error rates of less than  $10^{-9}$  are required.

Table 1  
HIPPI-to-ATM/SONET Converter Layer Functions

- Convergence Sublayer
  - Handling of Transmission Errors
  - Handling of Lost and Out-of-Sequence Cell transmissions
  - Flow Control
  - Timing and Source-Destination Synchronization
  - Segmentation and Reassembly
- Segmentation and Reassembly Sublayer
  - Segmentation of Messages into ATM Cells
  - Reassembly of ATM Cells into Messages
  - Detection of Lost and Out-of-Order ATM Cells
  - Detection of Bit Errors
- ATM Layer
  - Switching and Multiplexing
- SONET
  - Transmission, Reception and Conversion Between Electrical and Optical Signals

- **Error Control**

Forward error correction over the spacelink is provided by a shortened Reed-Solomon block code.

- **Data Structure**

Information is conveyed over the spacelink in the Asynchronous Transfer Mode (ATM), utilizing the framing, timing and overhead structures of the STS-3 Synchronous Transport Module. The latter consists of a rectangular array of 9 rows and 270 columns of octet-based information and is transmitted in 125 microseconds, for an effective transmission rate of 155.52 Mbit/s.

### **4.3 ACTS Earth Terminal Requirements**

To efficiently support the applications that are to be carried out over the distributed high performance computing environment the high data rate ACTS earth terminal is designed to meet the following communications requirements:

1. Support the transfer of up to three simultaneous full duplex continuous data streams at user data rates of approximately 150 Mbit/s in point-to-point mode between terrestrial interfaces at JPL and GSFC, utilizing the IF Switch Matrix in bent-pipe mode and hopping beams for uplink and downlink communications. Each data stream is transmitted in continuous mode without TDMA and over a separate transponder.

Transmission is in transparent pipe mode, with all protocol processing above layer 1 performed external to the terminals.

2. Provide a terrestrial interface capable of accepting and delivering up to 4 continuous optical carrier signals conforming to the OC-3 characteristics as given in ANSI specification T.106

- Line rate per optical carrier signal: 155.52 Mbit/s
- Required end-to-end bit error rate:  $10^{-10}$  or less

- Required Availability: 0.999
  - Optical Connector: As specified in ANSI/EIA 455 and ANSI/EIA 4750000-A
3. Perform the optical/electrical conversion between the spacelink and terrestrial signals. The conversion must conform to the following:
- Frame structure: STS-3 as defined in ANSI specification T.105, with nine columns of section and line overhead, one column of path overhead and 260 columns of payload
  - Frame synchronization: as specified in ANSI T.105
  - Baseband coding: CMI
4. Error control is carried out through forward error correction, utilizing a shortened Reed-Solomon Forward Error Correction Code with the following parameters:
- Block Length: 232 RS Symbols
  - Redundancy: 16 RS Symbols
  - Error Correction Capability: 8 RS Symbols
  - RS Symbol Size: 8 Bits
  - Number of RS Codewords per 1 ms Transmission Frame: 90
5. Carrier modulation employs binary phase shift keying at a modulation rate of 200 Mbit/s. The required signal-to-noise ratio  $E_b/N_0$  is approximately 8.9 dB.
6. The design of the spacelink protocol must take into consideration the requirements for framing and synchronization, sequence control, flow control and error control. The protocol specification includes:
- Modified TCP/IP to accommodate the high data rate and large round trip delay
  - Increased sequence number space to avoid sequence number wrap around
  - Increased window size through scaling

Protocol modifications are under development by Cray Research, and by the AURORA, CASA and NECTAR terrestrial gigabit testbeds.

#### **4.4 ACTS Earth Terminal Description**

A DARPA/NASA solicitation for high data rate earth stations (BAA92-12) is currently open, so the exact specifications of the terminal equipment used in the supercomputer network experiment will not be known until the solicitation is awarded. However, a good description of the terminal can be deduced from the BAA.

An ACTS high data rate terminal will consist of outdoor and indoor equipment. The outdoor antenna and RF communications assembly will be easily transportable and available in a "non impact" design configuration. The equipment will probably be assembled on a small trailer that can be towed by a van or light truck. The transportable equipment will be placed near the building housing the supercomputer, with the antenna able to "view" the ACTS satellite at its geostationary orbital location of 100 degrees west longitude. This results in the antenna pointing south-southwest from most CONUS locations. The physical size of the antenna is nominally in the 2.4 to 2.7 meter range for easy transportability and experimenter use. The low noise receiver is usually attached to the rear of the antenna with the remaining RF and baseband equipment rack mounted either in a small weatherproof house on the trailer, or located nearby in the computer room. The high power amplifier will produce about 200 watts of power, with effective radiated power being a function of station location vis-a-vis spot and scanning-beam arrangements of the satellite, along with rain fade operations and margins. The indoor HIPPI to AMT/SONET conversion equipment will be installed in a few standard sized racks, and will be located near the supercomputer. Nominal (115V) power lines and the communications cable(s) will connect the outdoor and indoor equipment.

The high data rate terminals to be used by the ACTS experimenters will be designed to be unmanned. Nevertheless, NASA is proposing to have Experiment Network Provider personnel under contract to install the equipment, provide interface services to the experimenters, and insure that the experiments are carried out in a timely and



efficient manner. It is essential that the experimenters, especially those with sophisticated interface requirements, be provided with these technical support services. The importance of the ACTS program along with the finite experiment and satellite time also require that trained technical personnel be available to support the experiments program.

## **5. ACTS Experiment Applications**

In this section we characterize several possible groups of applications for supercomputing networking between the GSFC and JPL Supercomputer Centers via the high data rate ACTS communications channels. The major emphasis in these applications is on a speed-up of the computationally intensive tasks. Of equal importance, however, is the potential of these experiments for enhancing the collaboration between the NASA centers in technical and scientific disciplines of common interest. Some examples of the latter are visible today, but many more are expected to appear as the ACTS system becomes reality, and the center project managers become aware of the availability of high data rate links.

In the following subsections we offer brief qualitative descriptions of several applications that suggest themselves in the near term for experimentation over the ACTS transmission channels.

### **5.1 Science Data Analysis Applications**

There are a number of science data analysis activities or experiments using the GSFC and JPL supercomputers which are facilitated by the near-gigabit data transfer rates proposed to be established in the context of the ACTS supercomputer networking experiment. Initially a series of science data analysis applications will be considered that will allow the size of the network and the sophistication level of the planned experiments to grow. In time more challenging applications can be developed involving higher data throughput and parallel rate instruments. These will eventually lead to experiments in support of the EOS Program.

Generally, the applications discussed below involve science data which either resides in the archives or is entering the NASA domain at GSFC or JPL and is reviewed or processed by experimenters using the large computing facilities at both GSFC and JPL at various stages of the analysis. The defining characteristic is the large volume of data

to be handled and distributed. The supercomputers themselves may be involved at one or both sites.

#### **5.1.1 Hubble Spacecraft Data Analysis**

One of the most important science data applications involves data analysis from the Hubble spacecraft. The instrument data, together with the ancillary and calibration data, resides with the Project at GSFC. The visualization tools and certain analytical expertise reside at JPL. Initial processing, image rectification and calibration, is performed by members of the investigative team at GSFC and the calibrated and registered data is transferred to JPL and other experimenters at current transfer rates. The volume of the refined data is likely to be somewhat less than that of the raw data sets. Visualization tools at JPL are used by the investigators there to mosaic and/or render the images. Final image products are delivered to collaborators at GSFC and elsewhere again via current communications circuits. The traditional or current best available transfer rates are T1 (1.544 Mbit/s). This requires long periods of data transfer which can be greatly expedited by the use of the ACTS link. In addition to a higher transfer rate, increased collaboration and analysis are expected during the ACTS experiment.

#### **5.1.2 Computational Fluid Dynamics Simulation**

The simulation of fluid dynamics processes using the Application Visualization System (AVS) software is an example of a satellite supercomputer networking science data application with two concurrent processes, one which produces the data stream and one which processes it. The simulation produces periodically a data "snapshot" and sends it to another process for analysis or visualization.

This simulation can be tested as part of an ACTS supercomputer networking experiment between the Goddard CRAY Y-MP and JPL CRAY Y-MP, using software supplied by AVS, Inc. This software is currently available for the CRAY machines and includes the facility for distributed computing within its environment. AVS is already

used as a visualization environment for the GSFC and JPL CRAY computers. Prior to the ACTS connection, the distributed computing test can be implemented using two local AVS machines, for example the Goddard CRAY and Convex, followed by test between the Goddard and JPL CRAYs over terrestrial circuits. Since the software already performs the necessary distributed computing functions, the implementation problems can be minimized and the test can be immediately applicable to any calculation that is structured to use the facilities of AVS.

### **5.1.3 Atmospheric Temperature Calculations**

The calculation of atmospheric temperatures at various altitudes from remotely sensed infrared data for each point on the earth's surface is another example of a satellite supercomputer networking science data application. This application requires completing many independent and computationally intensive tasks. The processing of temperatures for a point represents a parallel task since the data for each point is independent of the adjacent points. Parallel Virtual Machine (PVM) software will allow the implementation of such distributed applications across multiple Crays such as those at JPL and GSFC. Using this software, a process initiated on one machine assigns work to the other and accumulates the completed results from both itself and the other machine.

### **5.1.4 TOVS Pathfinder**

The proposed TOVS Pathfinder includes the TIROS Operational Vertical Sounder "TOVS", a polar-orbiting and satellite-borne instrument for measuring spectral radiance of the Earth and atmosphere at several wavelengths. With appropriate but difficult analysis, these data are useful for study of weather, climate, and geophysical processes. The TOVS system has flown on satellites since 1978 and corresponding instruments will continue in the EOS era. Analysis to date has emphasized weather prediction modeling. Guided by the TOVS Pathfinder Scientific Working Group and their Nov. 25, 1991 report to NASA and NOAA, it has been proposed that the archived data would be retrieved,

recalibrated, and made readily accessible for climate modeling and global change analysis. Currently the TOVS archived data accumulates at the rate of approximately 100 million bytes per day, an amount that is transported overnight via existing links. A small subset of the existing data has been transported to JPL in this fashion. Faster pathways are needed for handling the bulk of the data accumulated over a period of 12 years. Essentially all processing to date has been performed on the GSFC Supercomputer, with the tasks at JPL being interactive visualization and validation. The use of ACTS transport links with rates significantly higher than those in use today will permit enhanced interaction with the analysis process.

In the EOS era, the daily accumulation of TOVS-like data is expected to be many magnitudes above that of the past, thus representing an operational need for direct data linkages comparable to those available via an ACTS channel. Additionally, insofar as multiple collaborating investigators will be at multiple sites, concurrent delivery of newly calibrated data could become a requirement.

#### **5.1.5 Coupled Atmosphere-Oceans Climate Model**

A computationally intensive task which could potentially make full use of the ACTS communications channels between JPL and GSFC is a coupled atmosphere-oceans climate model. In concept, JPL researchers would lead the development of the oceans segment, while GSFC would provide the lead for the atmospheric part of the model, essentially paralleling the assignments for two of the Distributed Active Archive Centers (DAACs) of the EOSDIS. Time scales of the two models differ and conveniently support dissection of the problem into two coupled parallel processes. Close association with the data sets of the related DAAC would facilitate calibration to reality.

Mechoso et al. (Ref# 7) discuss decomposition and distribution of a climate model across multiple supercomputers linked by one of the regional terrestrial gigabit testbeds. These authors argue that with the gigabit interconnects, such decomposition and distribution will indeed enable the concurrent utilization of computing resources at geographically separated locations and speed up model execution. Numerical results are

presented for a particular decomposition of a UCLA climate model and parameters of interest, which give a very high CPU efficiency and lower bandwidth efficiency for the HIPPI (800 Mbit/s) link. For this case, the net average transfer rate is around 4-5 Mbit/s, although much higher peak rates would be needed if the communications and calculations do not completely run in parallel. The data rates to be available with the ACTS would comfortably handle this model, and would in fact permit much closer coupling of the segments if desired.

Mechoso et al. also briefly describe a preliminary version of the distributed application, wherein the oceans and atmospheres models are run as coupled processes at two widely separated supercomputer sites, using existing networking capabilities for communications. The net effective bandwidth obtained in competition with other users was about 180 kbit/s. Communications was not paralleled with computation and was dominant, consuming approximately two-thirds of the total clock time. This preliminary version showed the functional feasibility of distributing the application, but also showed it to be impractical, given the limitations of the current network. The application would appear to become practical with an effective bit rate above 10 Mbit/s, such as would be available with ACTS.

## **5.2 Animation and Visualization of Movie Frames**

The visual rendering of the many frames of a data product such as "Mars: the Movie" is a very computationally intensive task which can be executed on a number of computational platforms at JPL, including the SUN workstations, the CRAY, and others. The calculations are nearly independent from frame to frame. They are frequently executed in a distributed fashion using the "Explorer" software to control and coordinate the activity. In this, the basic user interface and control point is via one of the SUNs. None of the platforms can individually perform the rendering at anywhere near real time. Typical times are on the order of a minute per frame.

Initiating the rendering and animation process requires establishing the image data base on each of the platforms to be used, which for a typical Magellan Mosaic occupies

about 800 Megabytes of memory. Distributing this across the country via the ACTS would take only a minute or so, but would take over an hour at current internet rates. The rendered frames are typically 720 kilobytes at standard video resolution, and can be returned at a small fraction of the time taken to render them.

The basic computational work station for performing the rendering, namely the Sun-4, takes approximately 3-4 minutes per typical frame. If invoked, the CRAY is "just another workstation" in the system, albeit a quite powerful one. Because the rendering code consists of mostly scalar operations, it is not well matched to the CRAY, nor does it appear to be easily vectorizable. As a consequence, the CRAY-YMP (single processor) is only six times faster than the Sun. The code has been successfully ported to the Hypercube, and will be tried on the Delta. The new Solbourne performs at about 1 minute per frame, only a factor-of-two slower than the CRAY.

Even though the visualization rendering process itself is not a good use of the CRAY, the ACTS linkage could also facilitate the sharing of various other computing resources that are more appropriate for the visualization process. In addition, linkage between the two Centers at the near-gigabit data rates enabled by the ACTS would clearly facilitate the sharing of data and collaboration in data analysis.

### **5.3 Expansion of Potential Applications**

From this brief examination of the potential space science and visualization applications that can utilize the ACTS to link computing facilities between GSFC and JPL, we conclude that there are numerous activities already visible or underway that would benefit significantly from the availability of the ACTS communications capabilities. Both the overall number of potential applications and their individual requirements for data exchange can be expected to increase in the future, especially with the approach of the EOS era. We have made no attempt here to be exhaustive, nor has there been an intent to develop the precise specifications of the applications. Significantly more work is needed in this area to establish specific plans and proposals.

Taken individually, the applications described above bear the common thread that they could probably be implemented over communications circuits that offer a tenfold increase over current communications capabilities. They could, however, effectively use the hundred-fold increase in capacity provided by the ACTS communications channels. This additional capability facilitates human interaction with the data analysis process, thereby making it more flexible and productive. Applications which are truly dependent upon the near-gigabit rates of ACTS may not surface until some experience with intermediate data rate capabilities is accumulated.

In contemplating the use of ACTS in these applications we are led to a consideration of the many details in which the ACTS communications channels differ from those of the terrestrial-based regional gigabit networks. The main factor here is the greater latency resulting from the round-trip propagation delay of the satellite communications path. This delay will require changes in the error control protocols used, and is likely to induce changes in the strategy for distributing the applications elements.

As a final and important comment, we note that the specification of an application, and the corresponding experiment to be carried out over ACTS to support the application, require careful consideration of the following major elements:

- Statement of the experiment objectives
- Conceptual design of the experiment
- Specification of the performance parameters
- Specification of the hardware and software requirements for the computational resources
- Specification of the data communications and storage requirements
- Specification of the ACTS user access requirements
- Determination of the experiment time schedules
- Definition of the experiment management procedures
- Specification of the data collection and reporting requirements



## 6. Conclusions and Recommendations

The conclusions of this design study reflect the important dimensions of supercomputer networking for the ACTS program. Application experiments of the type described in this study have neither been performed in previous programs, nor can they be done efficiently over existing communications satellites. The importance of ACTS in this context is also being recognized by other federal research centers, universities, and industrial organizations that have supercomputers or high data rate requirements. Our team is being approached by a number of these entities with requests to join our proposed experiment or to receive reports or documentation that result from our experiment.

This study confirms that supercomputer networks using the ACTS satellite are feasible at bit rates up to 800 Mbit/s. Networking can be accomplished with either FDMA, TDMA or CDMA modulations. An examination of data rates concludes that the OC-3 rate of 155 Mbit/s can be implemented, and can be ready for tests and demonstrations shortly after the launch of the satellite. In addition, this rate is compatible with terrestrial networks that are planned to be available for testing during the same period.

The necessary earth terminal and interface equipment, described in this report, are presently not available. A joint DARPA-NASA procurement (BAA92-12) has been issued to acquire the terminal equipment, but it is not clear whether the terrestrial interface units (TIU) and the front-end processor (FEP) hardware and software are included in the procurement. Clearly this interface equipment is essential in order to connect a supercomputer to the ground terminal equipment.

It is concluded that while a number of excellent applications have been characterized, further emphasis and development must be given to the preparations required for the actual tests and demonstrations that are planned as part of the ACTS supercomputer network. These preparations need to be initiated prior to the launch of the satellite, and requirements made for optimization of the time available on the satellite for experimentation.

It is recommended that the importance of supercomputer networking via the ACTS program be recognized, with terminal equipment and experiment opportunities provided not only to the GSFC and JPL effort, but to other NASA centers such as LeRC, ARC, and LaRC, along with other U.S. government research centers such as the Naval Research Laboratory, the DOE National Laboratories, and the National Center for Atmospheric Research (NCAR). In addition, the University of Hawaii has established a budget for the procurement of a high data rate terminal and desires to be a part of the ACTS experimental program.

It is also recommended that priority be given to procure the necessary terminal and interface equipment for conducting the type of distributed high performance computing application experiments described in this study. If the interface equipment is not included in BAA92-12, a special effort must be made to complete the specifications and procure the TIU and FEP hardware and software.

It is recommended that the actual experimentation be recognized as an unusual opportunity to satisfy Earth and space science projects by NASA Centers, where these projects have requirements for collaboration of more than one Center, and where the implementation includes computationally intensive tasks that can be paralleled or coupled. This requirement will become very necessary in the EOS and EOSDIS programs.

Finally, it is recommended that the capabilities of ACTS for DHPC be demonstrated to and tested by the national long distance and regional carriers and the inter-urban Teleport networks. These entities should be invited to participate in ACTS DHPC experiments.

## 7. List of Investigators and Contributors

This study has been prepared with the dedication of and participation by a number of individuals. The many team members who worked on this study have also participated in the submission of an application to the ACTS Experiments Program Office (Code C) for this "Supercomputer Networking Applications" project to become an official ACTS experiment. As shown by the list of applications in section 6, the team members plan to conduct a series of tests and demonstrations over the ACTS system. The team is also aware of other research centers and institutions that have indicated an interest in being a part of this exciting DHPC project. Therefore, it is anticipated that our team will expand as the project matures.

The following list covers only the primary participants of this study. Certainly there are many other contributors and supporters that are too numerous to mention.

<u>Individual</u>	<u>Task</u>	<u>Organization</u>
B. Edelson	Principal Investigator	GWU
H. Helgert	Co-PI	GWU
N. Helm	Project Coordinator	GWU
T. Yarling	Project Assistant	GWU
P. Chitre	Co-PI	Comsat
W. Cook	Contributor	Comsat
W. Schmidt	Contributor	Comsat
G. Kirschner	Co-PI	Cray
R. Haley	Contributor	Cray
M. Halem	Co-PI	GSFC
D. Spicer	Contributor	GSFC
P. Gary	Contributor	GSFC
A. Zygielbaum	Co-PI	JPL
L. Bergman	Contributor	JPL
J. Layland	Contributor	JPL

## 8. Bibliography

Brandon, William T., Costello, Walter J., Supercomputer Network: An Operational Capability Using ACTS, AIAA 14th International Communications Satellite Systems Conference, Washington, D.C., March 1992.

Committee on Physical, Mathematical, and Engineering Sciences, Grand Challenges: High Performance Computing and Communications, The FY 1992 U.S. Research and Development Program.

Halem, Milton and Spicer, Daniel, Earth & Space Science Applications of Supercomputers, AIAA 14th International Communications Satellite Systems Conference, Washington, D.C., March 1992.

Helgert, Hermann, Networking Supercomputers Via Satellite Channels, Paper #235, AIAA 14th International Communications Satellite Systems Conference, Washington, D.C., March 1992.

Jacobson, V., Braden, R., TCP Extensions for Long-Delay Paths, RFC 1072, LBL and USC/Information Sciences Institute, October 1988.

Kirschner, Gerald, Supercomputer Networking Via Satellite, Paper #243, AIAA 14th International Communications Satellite Systems Conference, Washington, D.C., March 1992.

Mechoso, C., Chung-Chun, M., Farrara, J., Spahr, J., Distribution of a Climate Model Across High-Speed Networks, Supercomputing '91 Conference Proceedings, pp 253-259, November 1991.

Nicholson, A., Golio, J., Borman, D., Young, J., Roiger, W., High Speed Networking at Cray Research, Cray Research, Inc., Eagan, MN.